

Fast computing of scattering maps of nanostructures using graphical processing units

Vincent Favre-Nicolin,^{a,b,*} Johann Coraux,^c Marie-Ingrid Richard^d and Hubert Renevier^e

^aCEA-UJF, INAC, SP2M, Grenoble, France, ^bUniversité Joseph Fourier, Grenoble, France, ^cInstitut Néel, CNRS-UJF, Grenoble, France, ^dUniversité Aix-Marseille, IM2NP, UMR CNRS 6242, FST, Marseille France, and ^eLaboratoire des Matériaux et du Génie Physique, Grenoble INPMINATEC, Grenoble, France. Correspondence e-mail: Vincent.Favre-Nicolin@cea.fr

Scattering maps from strained or disordered nano-structures around a Bragg reflection can either be computed quickly using approximations and a (Fast) Fourier transform, or using individual atomic positions. In this article we show that it is possible to compute up to 4.10^{10} reflections · atoms · s⁻¹ using a single graphics card, and we evaluate how this speed depends on number of atoms and points in reciprocal space. An open-source software library (PyNX) allowing easy scattering computations (including grazing incidence conditions) in the Python language is described, with examples of scattering from non-ideal nanostructures.

© 2010 International Union of Crystallography
Printed in Singapore – all rights reserved

1. Introduction

Efficient computing of X-ray (and neutron) scattering from crystals has been the subject of intense work since computers became available. Except in the case of small structures (<1000 atoms) or small number of reflections (<1000), the method of choice has long been to use the Fast-Fourier Transform (Ten Eyck, 1973; Immirzi, 1973; Ten Eyck, 1977) of the crystal's scattering density. By computing this density inside the crystal's unit cell over a suitable grid (Ten Eyck, 1977; Langs, 2002), it is possible to compute structure factors at nodes of the reciprocal lattice.

In the case of strained (Takagi, 1969) or disordered (Butler & Welberry, 1992) crystals, the scattering must take into account a large part of the crystal (or possibly the entire crystal) instead of a single unit cell, in order to describe the departure from an infinite, triperiodic structure. This requires either using approximations, or a large computing power. Moreover, both strain and disorder lead to non-discrete scattering, so that the scattered amplitude must be evaluated on a fine grid around or between Bragg diffraction peaks. This type of computations can greatly benefit from fast calculations, which we will present in this paper.

This article is organized as follows: in section 2 we describe the formulas used for computing the scattering from an atomistic model, how it can be efficiently computed using a Graphical Processing Unit (GPU), and what performance can be achieved. In section 3 we present the open-source package PyNX which can be used to easily compute scattering with little programming knowledge. In section 4 a few examples are given.

2. Scattering computing from an atomistic model using a GPU

2.1. Theory

X-ray and neutron scattering can generally be calculated, in the kinematic approximation, as:

$$A(\mathbf{S}) = \int_V \rho(\mathbf{r}) \exp(2i\pi\mathbf{S} \cdot \mathbf{r}) dV = FT[\rho(\mathbf{r})] \quad (1)$$

where $A(\mathbf{S})$ is the scattered amplitude, \mathbf{S} is the scattering vector, $\rho(\mathbf{r})$ represents the scattering density (electronic or nuclear) at position \mathbf{r} inside the crystal, and FT denotes the Fourier transform. This equation can be used to determine the scattering from a crystal as long as $\rho(\mathbf{r})$ is described on a grid fine enough to resolve the atomic positions, which is easy if the crystal can be described from a single unit cell.

In the case of an aperiodic object (crystal with an inhomogeneous strain or disordered), it is simpler to compute the scattering from an atomistic model, which can be obtained using reverse Monte-Carlo (Nield *et al.*, 1995; Proffen & Welberry, 1997; Welberry & Proffen, 1998), atomic potentials combined with molecular dynamics or a direct minimization of the crystal energy (Keating, 1966; Stillinger & Weber, 1985; Tersoff, 1988; Plimpton, 1995; Niquet, 2006; Katcho *et al.*, 2009). The scattered amplitude is then derived from the atomic positions:

$$A(\mathbf{S}) = \sum_j f_j(S) \exp(2i\pi\mathbf{S} \cdot \mathbf{r}_j) \quad (2)$$

where $f_j(S)$ is the scattering length (either the Thomson scattering factor for X-rays or the nuclear scattering length for neutrons) of atom j and \mathbf{r}_j its position.

The number of floating-point operations (*flop*) required to evaluate equation (2) is approximately equal (see section 2.2) to:

$$N_{flop} \approx 8 \times N_{atoms} \times N_{reflections} \quad (3)$$

For a structure with $N_{atoms} = 8 \times 10^6$ (e.g. a cube of silicon of $54 \times 54 \times 54 \text{ nm}^3$) and $N_{refl} = 256 \times 256$ points in reciprocal space, this corresponds to 4.2×10^{12} *flop*, which can be

computer programs

compared to the current computing power of today's consumer micro-processors of $5 - 10 \times 10^9 \text{ flop} \cdot \text{s}^{-1}$ per core.

In the case of large nano (and micro)-structures, for which the description of the atomic structure is not possible in practice, a model based on continuum elasticity can be used, either with an analytical or numerical approach (see Stangl *et al.* (2004) and references within). The most popular method in the case of epitaxial nanostructures currently is the Finite Element Method-see Wintersberger *et al.* (2010) for a recent discussion. This model can then be used to calculate the scattering using groups of atoms:

$$A(\mathbf{S}) = \sum_j F_j(\mathbf{S}) \exp(2i\pi\mathbf{S} \cdot (\mathbf{r}_j^0 + \mathbf{u}_j)) \quad (4)$$

where $F_j(\mathbf{S})$ is the structure factor for the j^{th} block of atoms (generally a group of unit cells), \mathbf{r}_j^0 its original position, and \mathbf{u}_j its displacement from the ideal structure.

Assuming that all blocks of atoms are identical and on a triperiodic grid, it is possible to rewrite Eqn. 4 as:

$$A(\mathbf{S}) \approx F(\mathbf{S}) FT [\exp(2i\pi\mathbf{H} \cdot \mathbf{u}_j)] \quad (5)$$

where \mathbf{H} denotes the Bragg peak position around which the calculation is made, $F(\mathbf{S})$ is the structure factor calculated for a block of atoms, and $\mathbf{u}(\mathbf{r})$ the displacement field inside the crystal.

If the composition of the blocks of atoms vary (e.g. due to interdiffusion), it is also possible to include a variation of the average scattering density in the FT (Takagi, 1969):

$$A(\mathbf{S}) \approx F(\mathbf{S}) FT [\bar{\rho}(\mathbf{r}) \exp(2i\pi\mathbf{H} \cdot \mathbf{u}(\mathbf{r}))] \quad (6)$$

where $\bar{\rho}(\mathbf{r})$ is the relative scattering density in the crystal.

Both equation (5) and (6) allow the use of a *fast* Fourier transform, but are only valid as long as:¹

$$|(\mathbf{S} - \mathbf{H}) \cdot \mathbf{u}(\mathbf{r})| \ll 1 \quad (7)$$

Moreover, use of equations (5) and (6) with a *fast* Fourier transform restricts the computation of scattering on a triperiodic grid in reciprocal space - this is a limitation since modern data collection often use 2D detectors, and the measured points in reciprocal space are located on a *curved* surface (the projection of the detector on Ewald's sphere). Furthermore, as the resolution in reciprocal space is inversely proportional to the size in real space, analysis of high-resolution data using a FFT calculation demands a large model - even if the extent in reciprocal space is very limited.

Therefore, even if the speed of the FFT is optimal for large crystalline structures - for N points in real space, N points in reciprocal space are calculated with a cost proportional to $N \cdot \log(N)$ instead of N^2 - it is still interesting to consider a *direct* computation using equation (2) or (4) because it allows computation for:

- any assembly of points in reciprocal space
- from any structural model (no matter how severely distorted or disordered)

2.2. Implementation on a GPU

In order to achieve the calculations in a reasonable time, it is useful to consider current graphics cards as general-purpose GPU. This has already been reported in the scope of crystallography, for computing scattering maps from disordered crystals (Gutmann, 2010), powder pattern computing using the Debye equation (Gelasio *et al.*, 2010), and for single-particle electron microscopy (Schmeisser *et al.*, 2009).

To summarize basic principles behind GPU computing, it is possible to accelerate any calculation provided that:

1. it is **highly parallel**, *i.e.* the same formula must be applied on large amounts of data, independently
2. the number of **memory** transfers required is much smaller than the number of mathematical operations
3. the calculation pathway is determined in advance (at compilation time), which excludes any *if...then...else* operation in the inner computation loop

Moreover, many classical functions (e.g.: *exp*, *log*, *sqrt*, fused *sin - cos* evaluation,...) are highly optimized on GPUs - an algorithm requiring many such operations will be greatly accelerated.

Equation (2) fulfills all requirements, assuming that both the number of atoms and the number of points in reciprocal space are large ($\gg 1000$).




Figure 1

Computing speed depending on the number of atoms and reflections. These tests were run on a single nVidia GTX295 graphics card, using in parallel the two multiprocessors available on the card. N_{refl} indicates the number of reflections for the GPU calculations (black lines). The CPU (Central Processing Unit) curves (red lines) correspond to a computing using a vectorized (SSE-optimized) C++ code running on a single core of an Intel Core2 Quad Q9550 running at 2.83 GHz, for $N_{refl} = 10^2, 10^3, 10^4$ (the curves for 10^3 and 10^4 are almost identical).

The implementation presented in this article uses the CUDA (NVIDIA, 2010) toolkit. It is beyond the scope of this article to

¹ A more detailed study of this approximation will be presented in another article devoted to X-ray coherent diffraction imaging in Bragg condition

detail the exact algorithm used for computation, as the implementation is freely available as an open-source project (see section 3). However, it should be noted that the calculations are made in parallel for all reflections, with the atomic coordinates shared between parallel threads (to minimize memory transfers) - this method is optimal for large number of atoms. For some configurations (large number of reflections and small number of atoms), it may be more optimal to parallelize on the atoms and share the reflection coordinates between parallel processes.

The achieved speed is shown in Fig. 1, for a calculation of scattering for a random list of points in reciprocal space, and random coordinates for the atoms - the occupancy of all atoms is assumed in this test to be equal to 1, and the atomic scattering factor is not evaluated- in practice the atomic scattering factors can be factorized and represent a negligible amount of computing (see section 4.2) - the same is true for Debye-Waller factors.

As can be seen in Fig. 1, there is a strong dependence of the speed with the number of reflections and the number of atoms per second - the maximum speed ($\approx 46 \times 10^9 \text{ reflection} \cdot \text{atoms} \cdot \text{s}^{-1}$) is only reached if both numbers are larger than 10^4 .

Each couple (reflection, atom) corresponds to 8 floating-point operations (3 multiplications, 4 additions, one *sincos* evaluation)², so that the overall speed is equal to $\approx 367 \text{ Gflop} \cdot \text{s}^{-1}$. This can be compared to the peak theoretical speed of $1.7 \text{ Tflop} \cdot \text{s}^{-1}$ for this graphics card, which is only achieved when using fused add-multiply operations, without any bottleneck due to memory transfers.

By comparison, when computing on the CPU (see Fig. 1), the maximum speed is reached sooner: for 100 atoms and >1000 reflections, or for >1000 atoms when using >100 reflections. The top speed for a *single* core (Intel Core2 Quad Q9550 running at 2.83 GHz), using SSE-vectorized³ sine and cosine functions (Pommier, 2008), is $122 \times 10^6 \text{ reflection} \cdot \text{atoms} \cdot \text{s}^{-1}$ - ≈ 380 times slower than the GPU version. An *un-optimized* (without using SSE code) C++ code runs about 3 times slower, or ≈ 1000 times slower than the GPU version. Using multiple cores, the speed increases linearly (except for small number of atoms) with the number of cores.

2.3. Accuracy

As was already pointed out by Gutmann (2010), accuracy is an important issue since GPUs are most efficient when using single precision floating-point operations. Moreover, the accuracy of operations can be slightly relaxed (NVIDIA, 2010) compared to IEEE standards (IEEE, 2008).

For example, since single-precision floating-point use 24 bits mantissa (*i.e.* a relative accuracy of $\approx 10^{-7}$), precision may be expected to become problematic when the total scattered amplitude varies on more than 7 orders of magnitude.

A simple test can be made: computing the scattering for a linear, perfectly periodic chain of identical atoms, and comparing it to the analytical formula: $A(h) = \frac{\sin(\pi N \times H)}{\sin(\pi H)}$ (where H

is the reciprocal lattice unit and N the number of atoms in the chain). This is shown in Fig.2, for chains of atoms of different lengths. The discrepancy between the analytical calculation and the single-precision GPU calculation is clearly visible in the regions where the intensity is minimal - however in practice, the dynamic range where the calculation is reliable is always larger than 10^8 , and most of the time ($> 99\%$ of the points) around 10^{10} - these numbers refer to the intensity (squared modulus of the scattered amplitude).

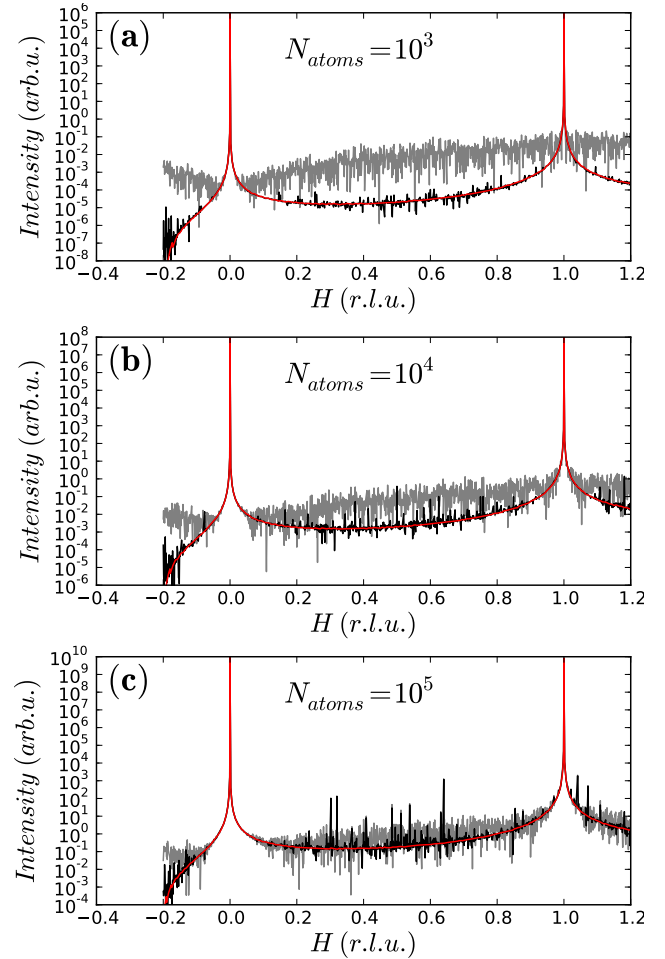


Figure 2

Comparison of the accuracy of the computation of the scattering between single-precision floating-point and an analytical model. The calculation is made for a linear chain of atoms with (a) 10^3 (b) 10^4 and (c) 10^5 atoms, using a perfectly periodic chain calculated using the GPU (black line), the analytical model (red line), and for a chain where the atoms are randomly displaced with a Gaussian distribution with a standard deviation of 10^{-3} (gray line). The atoms were located at $0, 1, \dots, (N_{atom} - 1)$, and the H coordinates are located every 0.001 in reciprocal lattice units (r.l.u.).

Such a dynamic range should be sufficient for most applications, as the practical range for experimentally measured intensities is usually lower, except in the case of perfect crystals. In

² Note that although the *sincos* operation is hardware-accelerated, it is 4 to 8 times slower than a simple addition. If the *sincos* evaluation is counted as 4 *flop*, the achieved speed is $\approx 500 \text{ Gflop} \cdot \text{s}^{-1}$.

³ SSE: Streaming SIMD Extensions ; SIMD: Single Instruction, Multiple Data

computer programs

Fig.2 a gray curve is superposed to the simulations, and corresponds to the GPU calculation for a chain of atoms with random displacements with a Gaussian standard deviation of 10^{-3} of their fractional coordinate. The error due to the single-precision computing is generally lower than the 'noise' level represented by the gray curve.

We have found that errors due to single precision floating point calculations were not significant *in practice*: indeed, most of the time structural models for which this type of computation is used are not ideal (see examples of simulated calculations using our code in Tardif *et al.* (2010) and Favre-Nicolin *et al.* (2010)) and therefore do not present a very large dynamic range (larger than 8 orders of magnitude).

It should however be noted that GPUs can also perform calculations using double precision floating point, but with a lower performance, as the number of available processing units are generally smaller (8 times in the case of CUDA graphics cards with capability less than 2.0 (NVIDIA, 2010)) than for single precision calculations. More recent graphics cards (available since mid-2010), using the Fermi architecture (http://www.nvidia.com/object/fermi_architecture.html) provide a higher computing power dedicated to double-precision computing (about half the speed of single-precision).

3. Open-source implementation: PyNX

Writing programs using GPU computing is a relatively complex process, as it is necessary to fine-tune the algorithm, notably in order to optimize memory transfers - which can make a very significant difference in terms of performance. For example an early version of the presented algorithm did not perform in a synchronized way between parallel computing threads, and its performance was $10\times$ slower than the final algorithm used. Moreover, all data has to be allocated both in the computer's main memory as well as on the graphics card, which can be tedious to write.

For this reason, we have written an open-source library, PyNX "Python tools for Nanostructure Xtallography", using the Python language. The main features of this software package are the following:

- Computing of scattering for a given list of atomic positions and points in reciprocal space does not require *any* GPU-computing knowledge
- it is possible to input either a list of (x, y, z) coordinates, or also include their occupancy (x, y, z, occ)
- the shape and order of the (x, y, z) and (h, k, l) coordinates (i.e. 1D, 2D or 3D, sorted or not) is irrelevant - all calculations are made *in fine* on 1D vectors
- the computation can be distributed on several GPUs - e.g. such cards as NVIDIA's GTX 295 are seen as two independent GPU units - the calculation is distributed transparently over the two GPU
- a pure SSE-optimized CPU computation is also available when no GPU is available, and can take advantage of all the computing cores available.

Three modules are available:

- `pynx.gpu`, which is the main module allowing fast, parallel computation of $\sum_j \exp(2i\pi \mathbf{S} \cdot \mathbf{r}_j)$, either using a GPU

or the CPU

- `pynx.fthomson`, which gives access to the analytic approximation for the X-ray atomic scattering factors (Prince, 2004) extracted from the cctbx library (Grosse-Kunstleve *et al.*, 2002)
- `pynx.gid`, which provides transmission and reflection coefficients at an interface (Dosch, 1992), which is required for grazing incidence diffraction analysis (Kegel *et al.*, 2000; Schmidbauer *et al.*, 2005; Richard *et al.*, 2009) using the Distorted Wave Born Approximation (DWBA), as is demonstrated in section 4.3.

4. Examples

4.1. Simple monoatomic cubic structure with $100 \times 100 \times 100$ unit cells




Figure 3

Example calculation of the 2D scattering of a strained cubic nanostructure (see text for details), around reflection (004). Intensities are color-coded on a logarithmic scale.

To compute the scattering around the (004) reflection of a simple cubic structure with a lateral size of 100 unit cells, the following code is used:

```
#Import libraries
from numpy import arange,float32,newaxis,log10,abs
from pynx import gpu

#Create array of 3D coordinates, 100x100x100 cells
x=arange(-50,50,dtype=float32)
y=arange(-50,50,dtype=float32)[: ,newaxis]
z=arange(0,100,dtype=float32)[: ,newaxis,newaxis]

#HKL coordinates as a 2D array
h=arange(-.1,.1,0.001)
k=0
l=arange(3.9,4.1,0.001)[: ,newaxis]

#The actual computation
fhkl,dt=gpu.Fhkl_thread(h,k,l,x,y,z,gpu_name="295")

#Display using matplotlib
from pylab import imshow
imshow(log10(abs(fhkl)**2),vmin=0,
       extent=(-.1,.1,3.9,4.1))
```

In this example, the calculation takes 0.93s on a GTX 295 graphics card. The library used for graphics display is matplotlib (<http://matplotlib.sourceforge.net/>).

Of course scattering from this cube could be calculated analytically - if we introduce a simple displacement field in the z-direction: $u_z = 10^{-6} \times z \times (x^2 + y^2)$, the following line can be inserted after the " z=arange..." instruction:

```
z=z+1e-6*z*(x**2+y**2)
```

The computed diffraction map is shown in Fig.3.

4.2. Bi-atomic structure from file

In the previous example, the atomic scattering factor is not taken into account - since this factor is the same for all atoms of the same type, it is easy to group all atoms of the same type together and calculate first $\sum_j \exp(2i\pi \mathbf{S} \cdot \mathbf{r}_j)$, and then multiply it by the value of the atomic scattering factor dependent on the position in reciprocal space (and the energy if anomalous scattering terms are to be taken into account), as well as the Debye-Waller factor. These atomic scattering factors can be extracted from the `pynx.fthomson` module.

Let us consider an InAs nano-structure, for which we have atomic coordinates in separate files `In.dat` and `As.dat`, each file having 3 columns corresponding to the x,y,z orthonormal coordinates (in nanometers). The scattering around reflection (004) for this data can be calculated in the following way (the f' and f'' resonant terms were taken manually from the cctbx library (Henke *et al.*, 1982; Grosse-Kunstleve *et al.*, 2002)):

```
#Import libraries
from numpy import arange,newaxis,sqrt,abs,loadtxt
from pynx import gpu,fthomson

#HKL coordinates as a 2D array
h=arange(-.1,.1,0.001)
k=0
l=arange(3.9,4.1,0.001)[: ,newaxis]

#Load orthonormal coordinates
xAs,yAs,zAs=loadtxt("As.dat",unpack=True)
xIn,yIn,zIn=loadtxt("In.dat",unpack=True)

#Convert to fractional coordinates
xAs/= .6036
yAs/= .6036
zAs/= .6036
xIn/= .6036
yIn/= .6036
zIn/= .6036

#Compute scattering
fhklIn,dt=gpu.Fhkl_thread(h,k,l,xIn,yIn,zIn,
                           gpu_name="295")
fhklAs,dt=gpu.Fhkl_thread(h,k,l,xAs,yAs,zAs,
                           gpu_name="295")

#Apply scattering factors at 10keV
s=6.036/sqrt(h**2+k**2+l**2)
fAs=fthomson.FThomson(s,"As")-1.64+0.70j
fIn=fthomson.FThomson(s,"In")+0.09+3.47j

#Full structure factor
fhkl=fhklAs*fAs + fhklIn*fIn
```

4.3. Grazing-Incidence Diffraction using the DWBA approximation

A specific module (`pynx.gid`) is available for grazing incidence diffraction - this module allows to compute the complex refraction index of a crystalline material (the substrate) and determine the reflection and transmission coefficients at the interface (Dosch, 1992). It is therefore possible to simulate grazing incidence X-ray scattering using the DWBA approximation, by taking into account the reflections before and after diffraction by the object at the surface, which influence the shape of the scattering in reciprocal space (Kegel *et al.*, 2000; Schmidbauer *et al.*, 2005; Richard *et al.*, 2009).

In Fig.4 is shown the simulated scattering for a germanium quantum dot on a silicon substrate. For the sake of demonstrating the use of PyNX, a simple analytical model is used:

- the quantum dot shape corresponds to a portion of a sphere, with a radius equal to 50 unit cells and a height of 20 unit cells.
- the germanium content varies linearly from 20% (bottom) to 80% (top)
- the relaxation (x,y,z being the fractional coordinates relative to the silicon substrate lattice) are:

$$\varepsilon_{xx} = \varepsilon_{yy} = 0.005 + z * 0.002 * (1 + \frac{1}{50} \sqrt{x^2 + y^2})$$

Moreover, in this example only the scattering from the quantum dot is calculated, ignoring any contribution from the substrate. The corresponding script is provided as a supplementary file. A more complete description of the scattering, taking into account scattering from the (strained) substrate (Schmidbauer *et al.*, 2005), could also be added in the future.




Figure 4

Simulated scattering from a germanium quantum dot on a silicon substrate, using the DWBA approximation. The calculation is made around the in-plane (400) reflection, and is plotted against H (reciprocal lattice unit) and the outgoing angle α_f . The location (outgoing angle) of the intensity maximum as a function of H varies as the in-plane strain changes with the height of the corresponding layer in the dot, due to interferences between the four scattered beams (Kegel *et al.*, 2000). See text for details.

5. Availability

The PyNX library is freely available from the project website at: <http://pynx.sourceforge.net>. It is open-source

computer programs

software, distributed under the CeCILL-B License (<http://www.cecill.info/index.en.html>), a permissive license similar to the FreeBSD license.

Although this library has been developed and tested only under linux, it should work on any operating system (including MacOS X and Windows) supported by the PyCUDA (Klößner *et al.*, 2009) library.

It has been tested on a variety of graphics card (9600 GT, GTX 295, 8800 GTX). Although it is recommended to use a dedicated graphics card (not used for display) for GPU computing, it is not a requirement - the library automatically cuts the number of atoms in order to decompose the calculation in batches which last less than 5s (a limit imposed by the CUDA library for graphics cards attached to a display). And it is also possible to use the CPU for calculations.

This library uses the scipy (<http://www.scipy.org>) and PyCUDA (Klößner *et al.*, 2009) libraries, and optionally the cctbx (Grosse-Kunstleve *et al.*, 2002) to determine the refraction index for the computing of transmission and reflection coefficients for grazing incidence X-ray scattering.

6. Conclusion

The main interest from this computing project is the ability to compute scattering for non-ideal structures without any approximation. This is particularly important for strained nanostructures where the calculation often uses a fast Fourier transform approximation, even though the displacements from the ideal structure are large. This could also be useful for coherent diffraction imaging in Bragg condition for strained nanostructures (Minkevich *et al.*, 2007; Robinson & Harder, 2009; Minkevich *et al.*, 2009; Favre-Nicolin *et al.*, 2010; Diaz *et al.*, 2010), especially in order to extend this method to severely distorted lattices (e.g. near an epitaxial interface with dislocations, a grain boundary,...).

A current limitation of this project is related to the toolkit used - the CUDA development package is the most popular GPU computing tool available at the moment, but it depends on a single manufacturer, and remains proprietary. An important development in that regard is the creation of the OpenCL language (<http://www.khronos.org/opencl/>), which is intended to allow GPU-computing *independently of the graphics card*. A future implementation of the proposed algorithm could use OpenCL and ensure its usability on a larger range of computing equipment.

The authors would like to thank Thierry Deutsch and Frédéric Lançon for helpful discussions during the development of this software package.

References

- Butler, B. D. & Welberry, T. R. (1992). *J. Appl. Cryst.* **25**(3), 391–399.
- Diaz, A., Chamard, V., Mocuta, C., Magalhães-Paniago, R., Stangl, J., Carbone, D., Metzger, T. H. & Bauer, G. (2010). *New J. Phys.* **12**(3), 035006.
- Dosch, H. (1992). *Critical Phenomena at Surfaces and Interfaces*. Springer Verlag, New York.
- Favre-Nicolin, V., Mastropietro, F., Eymery, J., Camacho, D., Niquet, Y. M., Borg, B. M., Messing, M. E., Wernersson, L., Algra, R. E., Bakkers, E. P. A. M., Metzger, T. H., Harder, R. & Robinson, I. K. (2010). *New J. Phys.* **12**(3), 035013.
- Gelisio, L., Ricardo, C. L. A., Leoni, M. & Scardi, P. (2010). *J. Appl. Cryst.* **43**(3), 647–653.
- Grosse-Kunstleve, R. W., Sauter, N. K., Moriarty, N. W. & Adams, P. D. (2002). *J. Appl. Cryst.* **35**(1), 126–136.
- Gutmann, M. J. (2010). *J. Appl. Cryst.* **43**(2), 250–255.
- Henke, B. L., Lee, P., Tanaka, T. J., Shimabukuro, R. L. & Fujikawa, B. K. (1982). *Atomic Data and Nuclear Data Tables*, **27**(1), 1–144.
- IEEE, (2008). IEEE standard for Floating-Point arithmetic. <http://dx.doi.org/10.1109/IEEESTD.2008.4610935>.
- Immirzi, A. (1973). *J. Appl. Cryst.* **6**(3), 246–249.
- Katcho, N. A., Richard, M. I., Landré, O., Tourbot, G., Proietti, M. G., Renevier, H., V. F., Daudin, B., Chen, G., Zhang, J. J. & Bauer, G. (2009). *Journal of Physics: Conference Series*, **190**, 12129.
- Keating, P. N. (1966). *Physical Review*, **145**(2), 637.
- Kegel, I., Metzger, T. H., Lorke, A., Peisl, J., Stangl, J., Bauer, G., García, J. M. & Petroff, P. M. (2000). *Phys. Rev. Lett.* **85**(8), 1694.
- Klößner, A., Pinto, N., Lee, Y., Catanzaro, B., Ivanov, P. & Fasih, A. (2009). *ArXiv*, **0911.3456**.
- Langs, D. A. (2002). *J. Appl. Cryst.* **35**(4), 505–505.
- Minkevich, A. A., Fohtung, E., Slobodskyy, T., Riotte, M., Grigoriev, D., Schmidbauer, M., Irvine, A. C., Novak, V., Holy, V. & Baumbach, T. (2009). *ArXiv*, **0909.4711**.
- Minkevich, A. A., Gailhanou, M., Micha, J., Charlet, B., Chamard, V. & Thomas, O. (2007). *Phys. Rev. B*, **76**(10), 104106–5.
- Nield, V. M., Keen, D. A. & McGreevy, R. L. (1995). *Acta Cryst. A Foundations of Crystallography*, **51**(5), 763–771.
- Niquet, Y. M. (2006). *Phys. Rev. B (Condensed Matter and Materials Physics)*, **74**(15), 155304–9.
- NVIDIA, (2010). CUDA, compute unified device architecture. http://www.nvidia.com/object_cuda_home_new.html.
- Plimpton, S. (1995). *J. Comput. Phys.* **117**(1), 1–19.
- Pommier, J., (2008). Simple SSE and SSE2 optimized sin, cos, log and exp. <http://grunthepeon.free.fr/ssmath/>.
- Prince, E. (2004). *International Tables for Crystallography, Mathematical, Physical and Chemical Tables*. Wiley, 3rd ed.
- Proffen, T. & Welberry, T. R. (1997). *Acta Cryst. A Foundations of Crystallography*, **53**(2), 202–216.
- Richard, M., Favre-Nicolin, V., Renaud, G., Schulli, T. U., Priester, C., Zhong, Z. & Metzger, T. (2009). *Appl. Phys. Lett.* **94**(1), 013112–3.
- Robinson, I. & Harder, R. (2009). *Nat. Mater.* **8**(4), 291–298.
- Schmeisser, M., Heisen, B. C., Luettich, M., Busche, B., Hauer, F., Koske, T., Knauber, K. & Stark, H. (2009). *Acta Cryst. D*, **65**(7), 659–671.
- Schmidbauer, M., Grigoriev, D., Hanke, M., Schäfer, P., Wiebach, T. & Köhler, R. (2005). *Phys. Rev. B*, **71**(11), 115324.
- Stangl, J., Holy, V. & Bauer, G. (2004). *Re. Mod. Phys.* **76**(3), 725.
- Stillinger, F. H. & Weber, T. A. (1985). *Phys. Rev. B*, **31**(8), 5262.
- Takagi, S. (1969). *J. Phys. Soc. Jpn.* **26**(5), 1239–1253.
- Tardif, S., Favre-Nicolin, V., Lançon, F., Arras, E., Jamet, M., Barski, A., Porret, C., Bayle-Guillemaud, P., Pochet, P., Devillers, T. & Rovezzi, M. (2010). *Phys. Rev. B*, **82**(10), 104101.
- Ten Eyck, L. F. (1973). *Acta Cryst. A*, **29**(2), 183–191.
- Ten Eyck, L. F. (1977). *Acta Cryst. A*, **33**(3), 486–492.
- Tersoff, J. (1988). *Phys. Rev. B*, **37**(12), 6991.
- Welberry, T. R. & Proffen, T. (1998). *J. Appl. Cryst.* **31**(3), 309–317.
- Wintersberger, E., Kriegner, D., Hrauda, N., Stangl, J. & Bauer, G. (2010). *J. Appl. Cryst.* **43**(6).